ICX-FIELDHAWK TECHNICAL PAPER
BNC-AN-105 REV A
AUGUST 2026

RF record and playback

Capturing a real radio environment so a device under test can face it again, identically, as many times as the investigation takes.

ABSTRACT

Some receiver failures only happen in the real world, and the real world is a poor test fixture: it does not repeat, cannot be scheduled, and will not hold still while one variable changes. Record and playback resolves that by turning the environment into a file. This paper sets out the complex baseband representation, the error terms that limit fidelity in a two-instrument capture-and-replay chain, the storage arithmetic that governs how long a scene can be, what synchronizing two generator channels actually requires, and how to prove a replay is faithful.

WHO THIS PAPER IS FOR

Receiver and system test engineers reproducing field failures on a bench. **GNSS, automotive and cellular validation teams** building repeatable scenario libraries. **EW and radar warning receiver engineers** who need range environments in the laboratory.

CONTENTS

1 Why the environment has to become a file

A receiver bug that appears once an hour at one intersection cannot be bisected against a radio environment that never recurs. The scene at any antenna is the superposition of every emitter in view, each with its own traffic, power control and mobility, and two successive observations of the same place differ in interferer timing, channel occupancy and fading state.

The obvious alternative fails the other way. A standards-compliant generated waveform repeats perfectly and contains nothing: no co-channel clutter, no adjacent-channel leakage, no unplanned transients. Receivers pass synthetic tests and fail in the field precisely because the field contains what the synthesis omitted.

Record and playback is the one method that is simultaneously real and repeatable at a chosen point in space. Capture the environment once as complex samples; replay it on the bench indefinitely, identically, on demand. The recording becomes a portable, versionable test asset: a scene that can be shared between teams, replayed against every firmware candidate, and attached to a bug report as the reproduction recipe.

TWO INSTRUMENTS, NOT ONE

The chain in this paper is deliberately split. A calibrated analyzer digitizes the air interface to a file; a separate vector signal generator modulates that file back onto a carrier. Neither instrument is the other, and their error terms do not combine the way a single-box datasheet implies. This paper keeps the two halves apart throughout, and names an instrument for either half only once the fidelity budget has said what each half has to do.

2 Complex baseband, and why bandwidth B needs sample rate B

A real band-limited signal at carrier frequency f_c can be written as an in-phase term on a cosine and a quadrature term on a negative sine. Those two real streams, taken together, carry both the amplitude and the phase of the modulation, and the generator reverses the operation exactly:

$$s(t) = I(t)\cos(2\pi f_c t) - Q(t)\sin(2\pi f_c t) \qquad (1)$$

where

I, Q is the in-phase and quadrature components of the complex envelope

f_c is the carrier frequency, in hertz

A real signal has a conjugate-symmetric spectrum, so each real sample carries one degree of freedom and unambiguous representation up to bandwidth B needs a sample rate of twice B. A complex sample carries two degrees of freedom, and a complex sequence at rate f_s represents a full span of width f_s about the tuned center without any symmetry requirement. So bandwidth B needs complex sample rate B rather than 2B. No information is gained: B complex samples per second is 2B real numbers per second either way.

PRACTICAL RULE

Usable bandwidth is always narrower than sample rate, because the anti-alias and decimation filters need a transition band. A ratio around 0.8 is typical, which is why quoting a file's sample rate as its bandwidth is a standard and avoidable error.

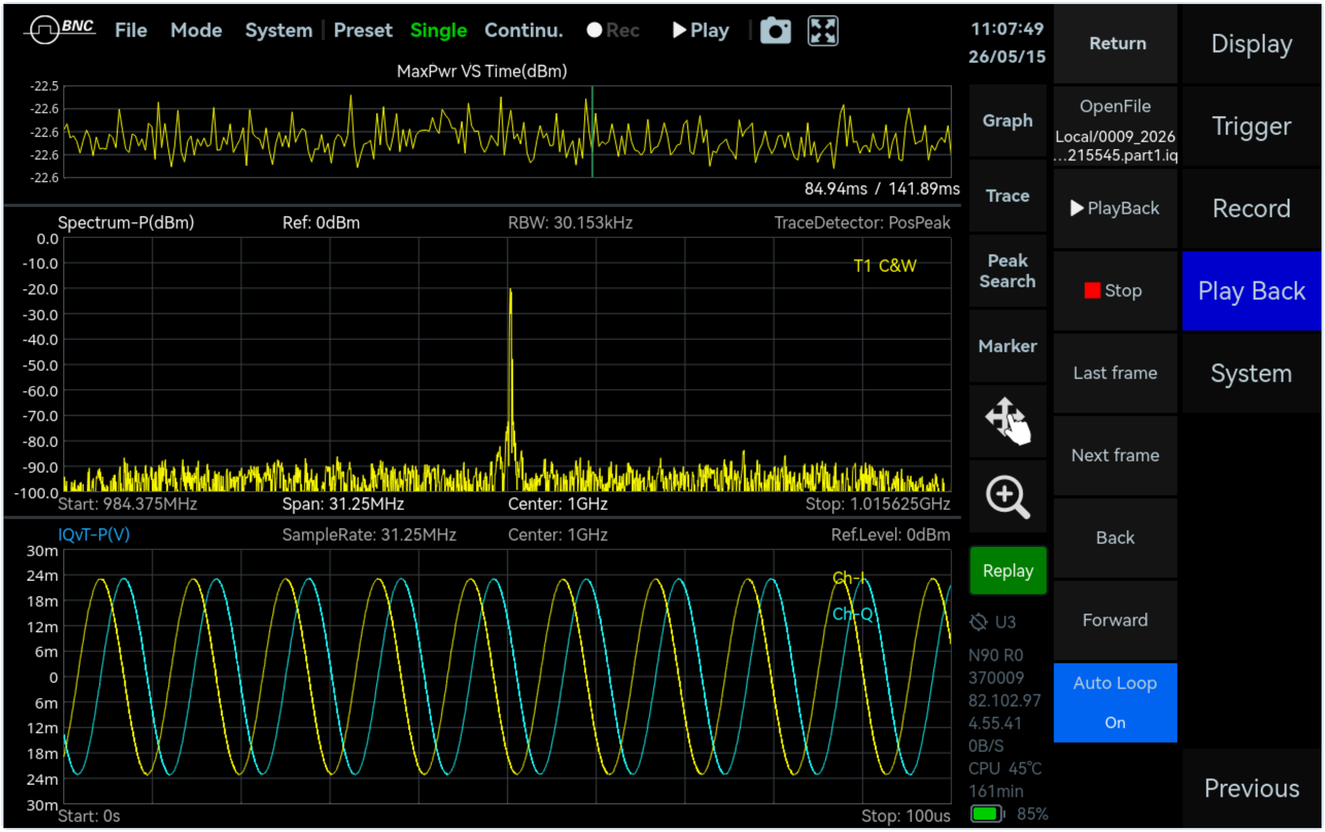

Figure 1. A recorded scene replayed on the ICX-FieldHawk: a stored IQ file at 1 GHz, 31.25 MHz of span at 31.25 MSPS, with I and Q a quarter cycle apart across 100 microseconds. Read the two numbers together. The span field is the sample rate, which is the error this section warns about, and 31.25 MSPS is 1.25 times the 25 MHz sustained figure, so the complex-Nyquist factor derived above is visible here as a setting. What matters before a scene is trusted is not the trace but the recorded state: center frequency, sample rate, reference level, decimation and an absolute timestamp. MEASURED

Figure 1 is that pair on a real instrument: the two traces are the same signal a quarter cycle apart, and the rate field underneath them is the ratio this section just derived, shown on a screen as a setting rather than in an equation.

3 What actually limits fidelity

Every impairment appears twice in a two-instrument chain, once in the analyzer and once in the generator, and where the two are uncorrelated they combine as a root sum of squares in error-vector terms. Images and spurs at related offsets do not obey that and can add coherently at the worst frequency in the scene. The two ends are not symmetric, and which end you spend money on follows from that.

> **THE SENTENCE TO REMEMBER**
>
> A replay can never be cleaner than its capture. The analyzer's noise, spurs, imbalance and phase noise are burned into the file, and to the generator they are indistinguishable from signal, so they are reproduced faithfully. The capture instrument sets the ceiling on scene fidelity. The generator only determines how little is lost from it.

Quantization, and the bits you actually get

Ideal quantization of a full-scale sine into b bits gives a signal to noise ratio of 6.02b plus 1.76 decibels. Real converters fall short, and the honest figure is the effective number of bits derived from measured signal to noise and distortion:

$$\mathrm{ENOB} = (\mathrm{SINAD} - 1.76) / 6.02 \tag{2}$$

where

$SINAD$ is the measured signal to noise and distortion ratio, in decibels

$ENOB$ is the effective number of bits

Two consequences follow immediately. Headroom costs bits: every decibel of unused input range below full scale is a decibel of signal to noise forfeited, so reference level at capture is a fidelity decision rather than a convenience, and a multi-emitter scene with a high peak to average ratio forces that headroom on you. And decimation buys it back, at ten times the log of the ratio between the original and final bandwidths: recording a 5 MHz emission out of a 100 MHz span is a factor of twenty, or 13.0 dB, which is more than two bits returned for the cost of a filter setting.

Quadrature imbalance, and why it is the dangerous one

Gain imbalance and quadrature phase error between the I and Q paths produce an image of the signal mirrored about the center frequency, with a rejection ratio of approximately:

$$\mathrm{IRR} \approx 4 / (\varepsilon^2 + \varphi^2) \quad (3)$$

where

ε is the fractional gain imbalance between the I and Q paths, $10^{\Delta G/20} - 1$ for a gain error of ΔG decibels, so 0.0058 at 0.05 dB

φ is the quadrature phase error, in radians

Table 1 shows how little imbalance it takes to cap image rejection.

Table 1. Image rejection against quadrature imbalance. Scaling both terms together costs 20 $\log_{10}$ of the factor, so 6 dB per doubling, and the two therefore enter the budget together with the worse of them dominating.

Gain imbalance	Phase error	Image rejection
0.05 dB	0.5 degrees	about 45.6 dB
0.1 dB	1 degree	about 39.6 dB
0.2 dB	2 degrees	about 33.6 dB
0.5 dB	5 degrees	about 25.6 dB
1.0 dB	10 degrees	about 19.5 dB

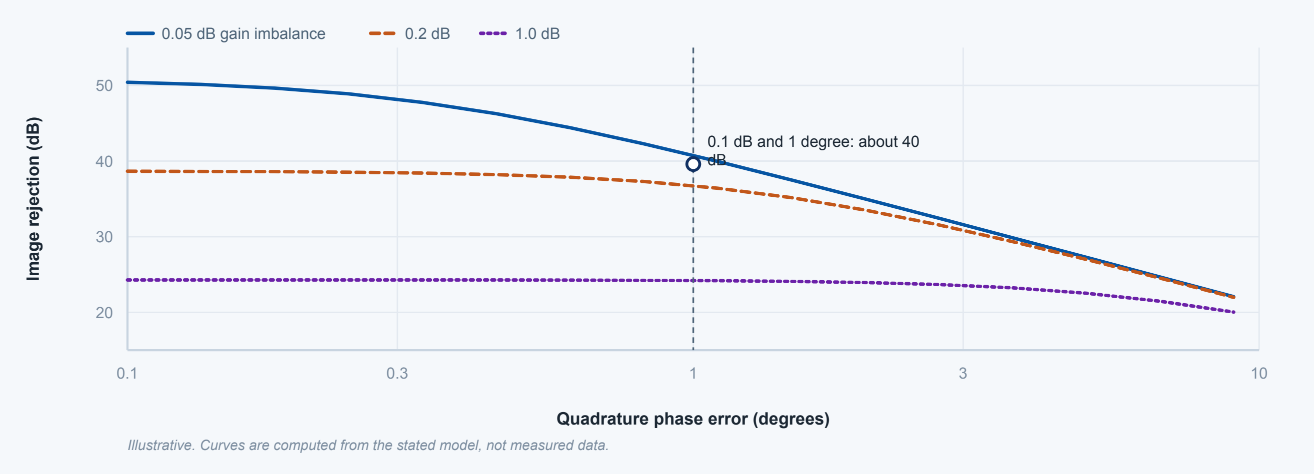

Figure 2. Tenths of a decibel and single degrees of imbalance cap image rejection near forty decibels, and both instruments in the chain contribute their own. Neither instrument's imbalance can be calibrated out of the other's, so the chain's rejection is set by whichever instrument is worse, not by the average of the two. DERIVED

Figure 2 plots the relation across the range that matters in practice, and the reading to take is the left-hand end. Each curve flattens onto a ceiling set by its gain imbalance alone, 50.8 dB at 0.05 dB and 24.3 dB at 1.0 dB, which no amount of phase correction gets past. That is why correcting one term and not the other buys so little.

Spectral asymmetry about a carrier is a genuine modulation feature: it is how single sideband differs from double, and it appears in asymmetric channel allocations. An imbalance image is spectral asymmetry manufactured by the instrument. In a replayed scene, the image of a strong emitter can land on top of a weak wanted signal mirrored about the center, and nothing downstream can tell the two apart. That is a corruption unique to record and playback, and it is worth checking for deliberately.

Reference offset and phase noise

A reference offset of x parts per million displaces the carrier by x times ten to the minus six times the carrier frequency, and it does so twice, in different ways. An offset at capture is frozen into the file as an apparent frequency error common to every emitter in the scene. An offset at replay shifts the reconstructed scene. Locking both instruments to a shared 10 MHz reference removes the second and can never remove the first. The same offset also stretches time: one part per million lengthens a sixty second recording by sixty microseconds and scales every modulation rate by the same factor.

4 How long a scene can be

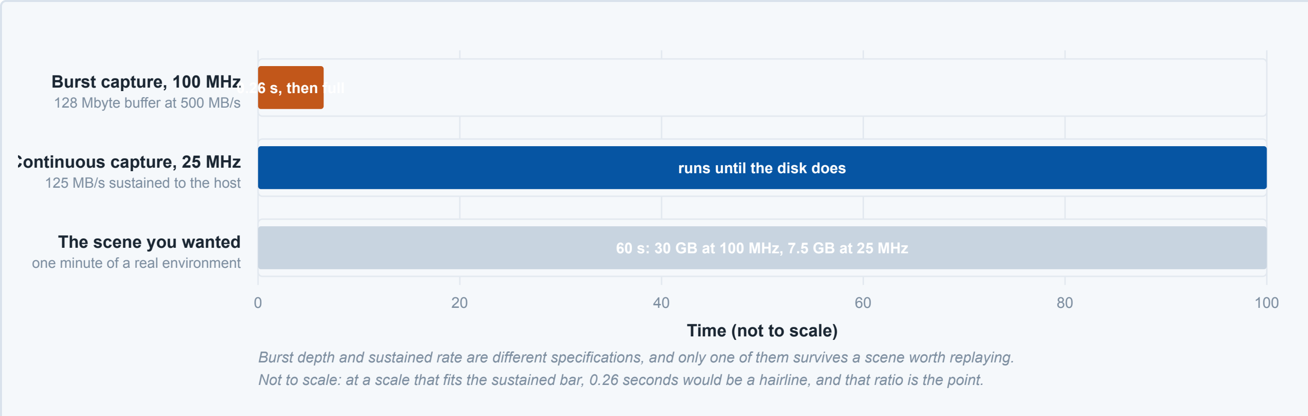

Figure 3. The two recording figures drawn on one timeline. A full-bandwidth burst buffer empties in 0.26 seconds, which is long enough to catch an event and far too short to be a scene; the sustained path runs indefinitely at a quarter of the bandwidth. A datasheet that quotes only the first number is quoting the one that flatters it. DERIVED

Figure 3 is the distinction a purchase specification has to force, and section 6 turns it into a numbered requirement. The arithmetic behind both bars is the same one line:

$$R = f_s \times 2 \times b/8 \tag{4}$$

where

R is the sustained data rate, in bytes per second

f_s is the complex sample rate

b is the bits per component, 16 in a conventional interleaved format

Table 2 turns that rate into scene duration.

Table 2. Storage for 16-bit complex capture at a sample rate 1.25 times the usable bandwidth. One hour of gap-free 100 MHz is roughly 1.8 terabytes, so RF recording is a storage engineering problem wearing an RF costume.

Usable bandwidth	Sample rate	Rate	One minute	One hour
1 MHz	1.25 MSPS	5 MB/s	300 MB	18 GB
10 MHz	12.5 MSPS	50 MB/s	3 GB	180 GB
25 MHz	31.25 MSPS	125 MB/s	7.5 GB	450 GB
100 MHz	125 MSPS	500 MB/s	30 GB	1.8 TB

This is why an honest datasheet quotes two different bandwidths. Burst capture writes into instrument memory at fabric speed and is bounded by capacity: the ICX-FieldHawk family publishes a 128 Mbyte buffer, which at 100 MHz and 16-bit components is about 0.26 seconds. Continuous capture has to sustain the slowest link in the chain indefinitely, and is published at 25 MHz. Filling a buffer and sustaining a transfer are different constraints.

Three levers manage the arithmetic. **Decimate at capture** and record the emission's bandwidth rather than the survey's, which is where the 13 dB of section 3 is actually spent. **Trigger** rather than stream, with enough pre-trigger depth to hold the leading edge, since the turn-on transient is exactly where emitter fingerprints live. **Choose the format deliberately**: 32-bit float doubles the file for no information gained over 16-bit integers, and lossless compression on near-noise IQ achieves very little.

> **HONEST LIMITATION**
>
> Lossy compression is available and should be ruled out explicitly for anything that will be replayed or used to train a classifier. It changes the data in precisely the dimensions that matter, and it does so invisibly.

5 Two channels, and what synchronized actually means

Multi-antenna receivers, wanted-plus-interferer scenarios and receiver blocking tests all need two generator outputs at once. The word synchronized covers three quite different requirements, and conflating them is the usual source of disappointment.

Table 3 separates the three rungs, in ascending cost and ascending capability.

Table 3. The three rungs of synchronization. A shared reference stops drift and says nothing about when playback starts; a common trigger aligns samples and says nothing about carrier phase. State the rung a scenario needs rather than specifying synchronized.

Rung	What it means	What it needs	Sufficient for
Baseband synchronous	no drift in rate between channels	one shared 10 MHz reference	nothing on its own
Time aligned	first sample out together, to the sample	a common trigger with deterministic latency, plus measured and nulled skew	throughput and interferer scenarios
Phase coherent	known and stable carrier phase between channels	shared or phase-locked local oscillators, phase-stable cabling, recalibration after every retune	angle of arrival and beamforming

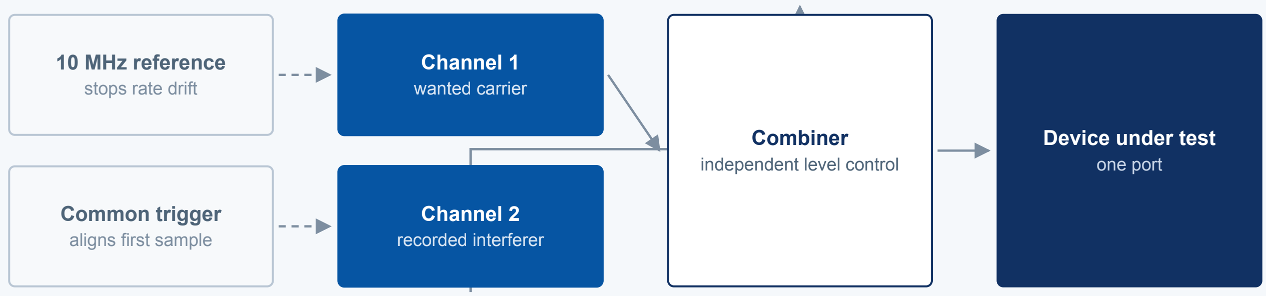

Splitting wanted and interferer across two channels unfreezes the ratio between them, which a single recorded scene cannot do.

Figure 4. The reference stops rate drift and the trigger aligns the first sample, and neither of them gives carrier phase. A blocking test needs only the middle rung; an angle-of-arrival test needs the top one, and a recalibration after every retune. SCHEMATIC

Figure 4 draws the arrangement. The reference and the trigger appear as separate lines because they solve separate problems, and a setup with one and not the other fails in a way that looks like a fault in the device under test.

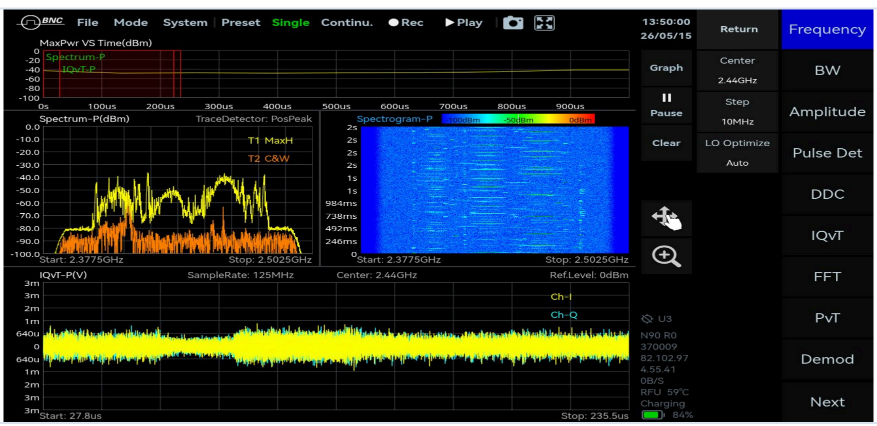

Figure 5. One acquisition of the 2.4 GHz band at 2.44 GHz, 125 MHz of span at 125 MSPS, again the sample rate rather than the usable 100 MHz, presented four ways at once: power against time, max-hold and clear-write spectra, a two-second spectrogram, and the I and Q traces across a 208 microsecond window. Every panel views the same samples, which is what recording an environment rather than a measurement of it means: the file still holds everything the analysis has not asked for yet. MEASURED

Figure 5 is worth holding against the fidelity budget above. Each of those panels is a different projection of one array of complex samples, so an impairment introduced at capture appears in all four and can be removed from none. That is the asymmetry of section 3, seen from the operator's side of the glass.

The wanted-plus-interferer case is worth naming because it solves a problem a single channel cannot. In a single recorded scene the ratio between the wanted signal and the interference is frozen at capture. Splitting them across two channels lets the operator sweep interferer power while the wanted carrier stays put, which is exactly what a blocking or desensitization test requires. Using a recorded real aggressor rather than a clean modulated blocker also reproduces mechanisms, reciprocal mixing, converter overload and gain-control pumping, that an idealized blocker understates.

6 Derived requirements

Sections 3 to 5 were a fidelity budget. Turned into requirements, they are what to put on a purchase specification for a record and playback bench, and Table 4 states each one with the term it comes from. Note which half of the chain each requirement lands on: that split is the whole reason this paper keeps the two instruments apart. R2, R3 and R4 are the ones a datasheet rarely answers directly, which is why section 8 measures them rather than reading them.

Table 4. Six requirements, and the third column is the one that saves money. R2 and R5 are answered by the analyzer alone, R6 by the generator alone, and only R1, R3 and R4 have to be met twice. A budget written without that column pays for the same decibel at both ends.

#	Requirement	Falls on	Derived from
R1	Complex sample rate at least 1.25 times the scene bandwidth	Both	Complex Nyquist, §2
R2	Effective number of bits stated as SINAD at the record bandwidth, not as converter bit depth	Capture	§3, quantization
R3	Image rejection stated, or gain and quadrature imbalance stated so it can be computed	Both, worse end dominates	§3, Eq. 3
R4	Phase noise at the offsets the modulation occupies, and a reference input	Both	§3, reference offset
R5	Sustained record depth at the scene bandwidth, distinguished from burst depth	Capture	§4
R6	Common reference and a hardware trigger, with carrier phase stated separately	Playback	§5, Table 3

Scoring the capture half

A requirements table nobody scores is a rubric, not an argument, so here is the capture half against the three of its own that a datasheet can answer. The ICX-FieldHawk family, driven by SpectraCore or by the published interfaces beneath it, publishes IQ sample rates to 125 MSPS with decimation from 1 to 4096, which satisfies R1 for any scene up to 100 MHz with the 1.25 factor in hand and gives the R5 lever directly. R5 itself is published as two separate figures rather than one flattering figure, which is the behavior the requirement was written to detect: burst recording to 100 MHz into 128 Mbyte, and continuous recording to a host at 25 MHz. R4 is answered in part: the reference input and the reference specification are published, TCXO below 1 ppm with an OCXO option below 0.15 ppm, and phase noise is published at a 10 kHz offset. One offset is not the offsets the modulation occupies, so R4 is a partial pass on any datasheet in this class, this one included.

Two of the six cannot be scored from any datasheet, ours or anyone else's, and saying so is the point of the exercise. **R2 is not published as SINAD at the record bandwidth by any vendor in this class**, so the effective bits behind a capture are a measurement the buyer has to make. **R3 is not published as image rejection either**, so the gain and quadrature imbalance that set it have to be measured with a single-sideband tone and Equation 3. Both take an afternoon on a bench that already exists, and both are worth more than a week of datasheet comparison, because they are the two terms that section 3 showed cannot be corrected out afterwards.

One application, carried through

Take a single job and run it down the table. A cellular device fails intermittently in one city and nowhere else, and the team cannot reproduce it. The scene is the 2.4 GHz band at a specific junction: an access point, a microwave oven, several handsets and a Bluetooth headset, all of them moving. R1 fixes the sample rate at 1.25 times the 83.5 MHz that holds the whole band, so 104 MSPS. Decimation runs in powers of two, so 104 is not a rate you can select: the choice is 125 MSPS or 62.5, and 62.5 carries only 50 MHz usable, so the whole band forces the full 125 and the file is a fifth larger than R1 alone implies. That is the ordinary way a requirement meets hardware. R5 then fixes the duration: at 500 MB/s a one-minute scene is 30 Gbyte, which the burst buffer cannot hold and the sustained path cannot reach at that bandwidth, so either the scene is decimated to the 25 MHz that actually carries the suspects, which is exactly the sustained figure R5 publishes and therefore buys unlimited duration, or it is cut to the seconds around the failure that the burst buffer holds. That is a decision about the experiment, forced by the R5 row of Table 4, and it is better made before the site visit than after. Decimation has one exception and this scene contains it: the headset hops across the whole 79 MHz, so a 25 MHz window keeps the access point and loses the hopper. For a frequency-hopping emitter the emission bandwidth *is* the survey bandwidth, and the only lever left is duration.

R6 decides the rest. Reproducing the failure only proves it reproduces; finding the threshold means sweeping the interferer while the wanted signal stays put, which a single recorded scene cannot do because the ratio was frozen at capture. Two channels, a common reference and a hardware trigger turn the recording into an experiment with a variable in it. **The scene answers what happened; the second channel answers how much margin was missing.** R2 and R3 apply to both, because an image or a lifted noise floor introduced at capture will be replayed as though the city had produced it.

7 Proving the replay is faithful

Fidelity is demonstrated, not assumed. Close the loop: replay the file into the same analyzer under SpectraCore, capture it again, and compare against the original. Because the same instrument bounds both captures, this measures what the generation and recapture chain added, which is a slightly pessimistic bound on the generator alone and an honest one.

1. **Lock both instruments to a common reference** before comparing, so frequency offset does not contaminate the result.
2. **Align coarsely by cross-correlation**, then finely by estimating and removing residual delay, phase offset and carrier frequency offset. Normalize power.
3. **Report a spectral comparison**: Welch-averaged power spectral densities overlaid, with maximum and root-mean-square deviation in decibels across the band. This catches passband tilt, images, added spurs and a lifted noise floor in one plot.
4. **Report error vector magnitude against the original file**, not against an ideal constellation. For record and playback the first is the honest metric, because the question is replay fidelity rather than signal quality.

Every step above is arithmetic on a complex array, which is worth saying plainly because it decides who can do it. The ICX-FieldHawk presents itself through SoapySDR, the vendor-neutral hardware abstraction layer, so the capture side of this loop can be driven from GNU Radio or from any other SoapySDR application, with the instrument calibration already applied to the samples. The cross-correlation, the residual-offset estimator and the Welch comparison are then existing blocks rather than a development project, and the amplitude reference underneath the fidelity number remains traceable. This path is demonstrated rather than specified; see the verification note.

Field evidence

One requirement recurs often enough to be worth naming: two generators under a single control port rather than two separately addressable instruments. It sounds like a packaging preference and it is not. Two channels that must be commanded separately are two instruments with two time bases, and section 5 is the reason a buyer who has run this loop before asks for one port. The synchronization rung a test needs is decided at purchase, not at setup, and it is the hardest thing to add afterwards.

8 Honest limits

- **Playback reproduces one point in space, not a channel.** The recording is the scene as seen by one antenna at one location, orientation and polarization. Replayed conducted into a device, the multipath geometry is frozen and moving the device under test changes nothing.
- **A single-antenna capture cannot reproduce a spatial channel.** One channel records one linear combination of the wavefield. A multi-antenna receiver fed copies of it sees a fully correlated channel, which is the one condition it is designed to exploit least.
- **Dynamic range is capped at capture.** A scene containing a strong blocker and a signal a hundred and forty decibels below it cannot be captured in one pass by any practical converter. That case needs the two-channel construction or separate captures.
- **Nothing outside the captured bandwidth exists.** The scene is brick-walled at the capture filter, so out-of-band blockers and image-frequency responses are untested by in-band replay.
- **Playback is deaf.** It cannot participate in closed loops: retransmission, power control, handover signaling, two-way ranging and transponder interrogation all vanish into a stimulus that never reacts. Record and playback tests receivers. It does not emulate networks.
- **Where other tools are right.** A channel emulator when the need is parametric, statistically controlled propagation; an over-the-air chamber when spatial behavior is under test; a scenario simulator when the trajectory must be authored rather than experienced. Complements, not rivals.

9 Summary, definitions and further reading

Table 5 is the block to lift into an evaluation matrix.

Table 5. Lift this whole. Two rows carry a caveat instead of a clean number, and those two are the ones to check against any vendor's datasheet before trusting the rest of it.

Quantity	Relation or value	Condition
Complex sampling	$f_s \geq B$, not 2B	usable BW about 0.8 f_s
Data rate	$f_s \times 2 \times b/8$	500 MB/s at 125 MSPS, 16-bit
Effective bits	(SINAD − 1.76) / 6.02	measured, not nominal
Image rejection	$4 / (\varepsilon^2 + \varphi^2)$	about 39.6 dB at 0.1 dB and 1 degree
Reference offset	$\Delta f = x \times 10^{-6} \times f_c$	also scales elapsed time
Reference stability	TCXO < 1 ppm; OCXO option 01 < 0.15 ppm	published as a frequency bound, not as holdover drift
Phase noise	−107.5 dBc/Hz (400); −101.6 (090)	1 GHz carrier, 10 kHz offset; one offset does not answer R4, which asks for the offsets the modulation occupies
Burst capture	128 Mbyte, about 0.26 s at 100 MHz	16-bit components
Continuous capture	25 MHz	sustained to host; the host and its storage have to keep up, so confirm the interface and the sustained write rate before relying on it
IQ sample rate	up to 125 MSPS, decimation 1 to 4096	powers of two

Capture values from the ICX-FieldHawk datasheets. Vector signal generation is a separate Berkeley Nucleonics instrument and is not an ICX capability.

Complex envelope
The in-phase and quadrature pair that together represent a modulated carrier at baseband, carrying both amplitude and phase.

Effective number of bits
The resolution a converter actually delivers, derived from measured signal to noise and distortion rather than from its nominal width.

Image rejection ratio
The ratio between a wanted signal and the mirror of it produced by quadrature imbalance.

Symbols used in this paper.

Symbol	Meaning	Units
f_c, B	Center frequency, signal bandwidth	Hz
f_s	Complex sample rate	samples per second
b	Bits per IQ component	bits
ε, φ	Gain imbalance, quadrature phase error	fraction, radians
IRR	Image rejection ratio	dB
SINAD, ENOB	Signal to noise and distortion; effective number of bits	dB, bits
R	Data rate	bytes per second

ENOB effective number of bits
GNSS global navigation satellite system
IQ in-phase and quadrature
IRR image rejection ratio
SINAD signal to noise and distortion ratio

Further reading

- *The Signal Metadata Format Specification*, the open standard for pairing an IQ file with the center frequency, sample rate, datatype and timestamp it needs to remain usable.
- ANSI/VITA 49.0 and 49.2, the packetized transport that carries IQ between instruments in real time with per-

packet timing and context.
- IEEE Std 1241, terminology and test methods for analog to digital converters, which defines the SINAD, ENOB and spurious-free dynamic range used here.
- B. Razavi, "Design considerations for direct-conversion receivers," *IEEE Transactions on Circuits and Systems II*, vol. 44, no. 6, 1997, for the imbalance and image relation.

If you are building a scenario library and want the fidelity budget worked against your own device under test, our application engineers would be glad to go through it with you, including the cases where a channel emulator rather than a replay is the right answer.

VERIFICATION NOTE

The following values in this paper are not yet confirmed against a published Berkeley Nucleonics datasheet and are marked verify in the text. They must be confirmed before this paper is released.

- Vector signal generation is not an ICX-FieldHawk capability. The only published generation feature in the family is an optional built-in signal generator on the USB models. Replay requires a separate Berkeley Nucleonics vector signal generator, and the specific model and its bandwidth should be confirmed before quoting a scene bandwidth.
- IQ file and bit format details, including interleaving order and scale convention, are not currently published and should be confirmed with engineering.
- The SoapySDR and GNU Radio path in section 6 is documented in a platform integration guide and demonstrated with captured results, but does not yet appear on a datasheet. The distribution route for the driver package is not settled, so no package name, repository or driver string is printed here.
- The purchasing pattern described in the field evidence section is reported from account history rather than a published source, and no end user, vendor or country is named.

Berkeley Nucleonics Corp.
Test, Measurement and Nuclear Instrumentation since 1963
For details: info@berkeleynucleonics.com or 800-234-7858

BNC-AN-105 Rev A
August 2026